arXiv:1405.3848v2 [math.GT] 3 Feb 2015

Sphere Recognition: Heuristics and Examples

Michael Joswig* Frank H. Lutz†
Mimi Tsuruga‡

Institut für Mathematik, TU Berlin, Straße des 17. Juni 136, 10623 Berlin, Germany

Abstract

Heuristic techniques for recognizing PL spheres using the topological software `polymake` are presented. These methods have been successful very often despite sphere recognition being known to be hard (for dimensions $d \geq 3$) or even undecidable (for $d \geq 5$). A deeper look into the simplicial complexes for which the heuristics failed uncovered a trove of examples having interesting topological and combinatorial properties.

Keywords: PL manifolds; discrete Morse theory; computational topology

1 Introduction

The sphere recognition problem often arises in the guise of *manifold recognition*; that is, deciding whether a given (finite abstract) simplicial complex triangulates some manifold then determining its type. In the piecewise linear (PL) category, recognizing whether a given complex triangulates a PL manifold can be reduced to *PL sphere recognition* since the links of all vertices of the given complex need to be PL spheres. The following is a (very incomplete) list of scenarios where manifold recognition can be used:

*Partially supported by the Einstein Foundation Berlin and DFG within the Priority Program 1489.

†Partially supported by VILLUM FONDEN through the Experimental Mathematics Network and by the Danish National Research Foundation (DNRF) through the Centre for Symmetry and Deformation.

‡Supported by the Berlin Mathematical School (BMS).

- ▷ Enumeration. When enumerating triangulations of manifolds of a given dimension with a fixed number of vertices or facets, we want to ensure that the objects produced are indeed manifolds and discard all others [13, 15, 53].

- ▷ Topological Constructions. Various topological manifold constructions can be discretized so that the objects of interest can be studied with the help of a computer. To ensure the discretization has been carried out correctly, we want to confirm the manifold property. In practice, this test effectively detects the majority of construction errors [2, 52, 55].

- ▷ Meshing. The goal here is to obtain a triangulation of a hypersurface in some (higher-dimensional) Euclidean space by sampling; see, e.g., [47]. As in the case of the topological constructions, we want to verify that the triangulation is non-degenerate.

Here we will only consider closed manifolds encoded as finite abstract simplicial complexes. However, our methods can be easily modified to deal with manifolds with boundary or more general cell complexes. For instance, the second barycentric subdivision of a pseudo-simplicial complex (having simplices as faces and allowing identifications on the boundaries) is a simplicial complex.

Before we describe our approach, we briefly look at the history of the sphere recognition problem. While sphere recognition is trivial in dimensions $d \leq 2$, Sergey P. Novikov showed that the problem is undecidable when $d \geq 5$; his proof can be found in the paper [56], see also [19]. Since then most of the research focused on $d = 3$. Rubinstein [46] and Thompson [54] proved that 3-sphere recognition is decidable. More recently, Schleimer [48] showed that 3-sphere recognition lies in the complexity class NP. Hass and Kuperberg [33] announced that this problem also lies in co-NP, provided that the Generalized Riemann Hypothesis holds. The problem of 4-sphere recognition is still open. It is even unclear whether there may or may not be *exotic* 4-spheres which are homeomorphic, but not PL homeomorphic to the standard 4-sphere $\mathbb{S}^4$.

Our main application is to certify whether a given simplicial complex is a combinatorial manifold. We do this by showing that all face links are PL-spheres. To this end our guiding principle will be the following key result.

Theorem 1 (Whitehead [57], Forman [26, 27]) *A combinatorial d-manifold is PL-homeomorphic to the standard PL d-sphere $\mathbb{S}^d$ if and only if it has some subdivision which admits a* spherical *discrete Morse function, i.e., a discrete Morse function with exactly one critical* 0*-cell and exactly one critical d-cell.*

A discrete Morse function may be encoded as an acyclic partial matching in the Hasse diagram [18]. The *critical* faces are those which are unmatched. Our strategy for sphere recognition is to combine the search for a spherical discrete Morse function with the computation of homology and fundamental groups and bistellar flips. This is explained in detail in Section 2. Section 3 will report on the relevant data structures and their implementation in `polymake` [5, 31].

Our procedure must overcome a number of major difficulties which include the following.

- ▷ Deciding whether a discrete Morse function with at most a fixed number of critical cells exists is NP-hard [37, 41].
- ▷ There are combinatorial d-spheres that do not admit any spherical discrete Morse function [9, 11].
- ▷ Adiprasito recently showed that a sufficiently large iterated barycentric subdivision of any PL sphere is polytopal [1] (and thus admits a spherical discrete Morse function). However, as pointed out above, for $d \geq 5$, recognizing the PL d-sphere is undecidable. This implies that a priori there is no bound on the number of barycentric subdivisions required to admit a spherical discrete Morse function.
- ▷ In iterated barycentric subdivisions, finding a spherical discrete Morse function quickly becomes cumbersome [2].

In Section 4, we will demonstrate that despite these drawbacks, finding optimal discrete Morse functions (within some 'horizon') is often surprisingly easy, even for large input; see also [2, 10]. Not surprising, however, is that there are also other kinds of input for which our methods fail. A thorough analysis of the reasons for failure reveals several interesting families of simplicial complexes. In this sense our sphere recognition procedure can be used in two ways: either as an effective method or as a device to systematically produce instructive examples. For instance, we obtain many non-collapsible simplicial complexes which are contractible.

We will work in the PL category. But for completeness, we collect here some known results about the recognition of non-PL spheres. A *proper homology sphere* is a topological manifold with the same homology as the sphere of the same dimension, but which is not homeomorphic to the sphere. By the double suspension theorem of Edwards [25] and Cannon [16], the double suspension of every proper homology d-sphere Σ^d is a true (d+2)-dimensional sphere. For any triangulation of Σ^d the double suspension $\mathrm{susp}(\mathrm{susp}(\Sigma^d))$ is a non-PL triangulation of $\mathbb{S}^{d+2}$ (since we have Σ^d sitting in its double suspension as the link of an edge). One way to recognize such non-PL triangulations would be to first reduce a given complex in size by using bistellar flips and then comparing the resulting triangulation with, say, reference triangulations of double suspensions for particular homology spheres — but this will have a limited chance of success. Surprisingly, there are non-PL balls and non-PL spheres with perfect discrete Morse vectors, see [2, 7]. This does not contradict Theorem 1 as some face links will fail to have perfect discrete Morse vectors.

This text is the full version of the extended abstract [36].

2 An Inductive Sphere Recognition Procedure

Throughout this text, let K be a d-dimensional (finite abstract simplicial) complex with n vertices and m facets. A *facet* is a face that is maximal with respect to inclusion. A d-dimensional complex is *pure* if each facet has exactly $d+1$ vertices. A codimension-1-face in a pure complex is called a *ridge*.

To verify whether K is a PL d-sphere, there are three elementary combinatorial checks that are useful to perform first. These checks are fast; their running time is bounded by a low-degree polynomial in the parameters d, m and n. If one of the checks fails, this will serve as the certificate that K is not a sphere.

(1) Check if K is pure.

(2) Check if each ridge is contained in exactly two facets.

Success in these two tests will ascertain that K is a *pseudo-manifold* (without boundary). A pseudo-manifold K of dimension $d = 0$ is the 0-dimensional sphere $\mathbb{S}^0$; it consists of two isolated vertices.

(3) If $d \geq 1$, check if the 1-skeleton of K is a connected graph.

A connected pseudo-manifold K of dimension $d = 1$ is a polygon, and thus triangulates the 1-dimensional sphere $\mathbb{S}^1$.

The pseudo-manifold property of a simplicial complex is inherited by all face links. In particular, a connected pseudo-manifold of dimension 2 is a triangulation of a closed surface or of a closed surface with pinch points. A *pinch point* has multiple disjoint cycles as its vertex link.

A d-dimensional pseudo-manifold is a *combinatorial d-manifold* if all vertex links are PL homeomorphic to the boundary of the d-simplex. In particular, a combinatorial d-manifold is a triangulation of a PL d-manifold. This recursive nesting of PL structures suggests an inductive check of the face links of K by dimension starting with 0-dimensional links and proceeding up. In this way the Steps (1), (2) and (3) serve as the base case for our inductive procedure.

A (connected) 2-dimensional pseudo-manifold K whose vertex links are single cycles is a combinatorial 2-manifold and triangulates a closed surface. If, in addition, the Euler characteristic of K is 2, then K is $\mathbb{S}^2$.

The sphere recognition problem becomes more interesting for $d \geq 3$. Despite the fact that exact methods for $d = 3$ exist [46, 54], employing a heuristic approach even in the 3-dimensional case often turns out to be efficient.

We begin by computing the *Hasse diagram* of the complex K; this is the directed graph having one node per face and a directed edge for each pair of incident faces whose dimensions differ by one. The orientation of the edges is merely a matter of convention; here we direct the edges towards the higher-dimensional faces. A proper implementation of this step along with a careful design of the relevant data structures is crucial for our heuristic approach; see Section 3.1 below for details.

(4) Compute random discrete Morse functions and check if any are spherical [10].

After the initial tests (1), (2), and (3) we need to verify that K is, indeed, a combinatorial manifold. A pure d-dimensional simplicial complex K is a combinatorial d-manifold if and only if for any proper i-face F of K, where $0 \leq i \leq d-1$, the link of F in K is a PL $(d-i-1)$-sphere. Notice that if all links of i-faces are PL spheres, then all links of $(i-1)$-faces are combinatorial $(d-i)$-manifolds, which is a necessary condition to verify that the links of $(i-1)$-faces are PL spheres. In this way, this property is recursive. In practice, however, a recursive method is likely to encounter repetitions so it is preferred to work inductively in a level-by-level approach.

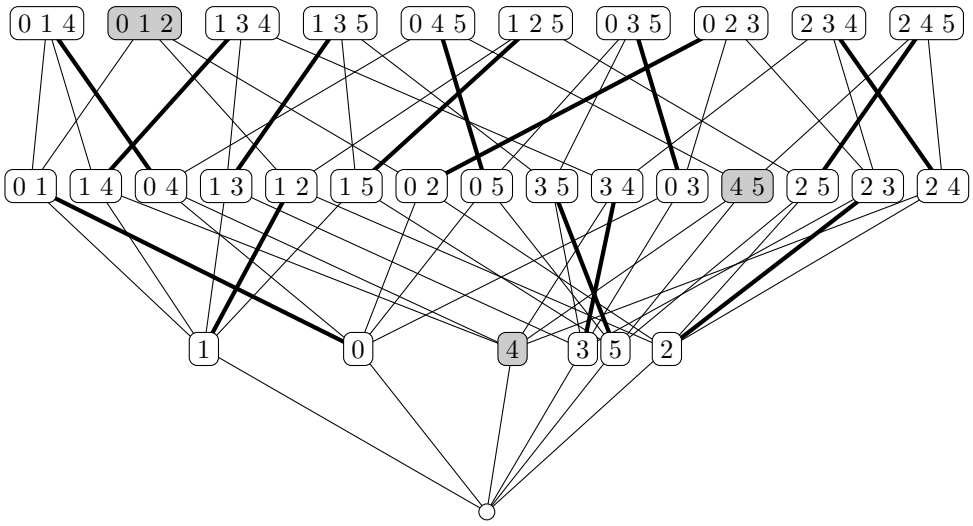

Figure 1: Acyclic matching in the Hasse diagram of $\mathbb{RP}^2_6$. The three unmatched cells are marked.

According to Theorem 1, any combinatorial d-manifold that becomes collapsible after the removal of one facet is a PL d-sphere. This statement is equivalent to the existence of an *acyclic matching* in the Hasse diagram with exactly two *critical cells*. That is, a matching given by the pairings induced by the performed elementary collapses such that if the edges of the matching are reversed, the resulting directed graph is acyclic and has precisely two unmatched nodes: one representing a facet and one representing a vertex. Chari [18] showed that this translates into the language of discrete Morse theory, as developed by Forman [26, 27], as follows: to every acyclic matching in the Hasse diagram there is a corresponding discrete Morse function with the same number of critical cells. For a general discrete Morse function on a d-dimensional simplicial complex, the *discrete Morse vector* $(c_0, c_1, \ldots, c_d)$ counts the critical cells per dimension. A discrete Morse function is *spherical* if its discrete Morse vector reads $(1, 0, 0, \ldots, 0, 1)$. Let $\mathbb{F}$ be some field. A discrete Morse vector is $\mathbb{F}$-*perfect* for K if $c_i = \beta_i(K; \mathbb{F})$ for $0 \leq i \leq d$. If K is a sphere or is contractible, then this definition is independent of the field and we simply say their respective discrete Morse vectors are *perfect*. See Figure 1 for the Hasse diagram of a triangulation of the real projective plane $\mathbb{RP}^2$ on six vertices; an acyclic matching with three critical cells is marked. The corresponding discrete Morse vector $(1, 1, 1)$ is $\mathbb{Z}_2$-perfect.

A randomized search for small discrete Morse vectors was introduced in [10]. This approach traverses the Hasse diagram level-wise from top to bottom. The free faces for elementary collapses are chosen at random; if there are no free faces, a face of the current maximal dimension is chosen at random, marked critical, and removed. For a discussion on how we randomize, see Section 3.3 below.

Brehm and Kühnel [13] already used a basic version of (4) to show that some 8-dimensional simplicial complex with 15 vertices is a combinatorial 8-manifold.

(5) Compute the homology and check if it is spherical.

Computing the simplicial homology modules is a fairly standard procedure. For field coefficients this reduces to Gauss elimination being applied to the (simplicial) boundary matrices. Over the integers the computation of the Smith normal forms of the boundary matrices serves as a replacement; see [42, §11]. Usually, we compute with integer coefficients. A necessary condition for K to be a sphere (PL or not) is $H_d(K) \cong \mathbb{Z}$ and all other (reduced) homology groups vanish. In this case we say that K has *spherical homology*. For more details, see Section 3.2 below.

(6) Compute a finite presentation of the fundamental group $\pi_1(K)$ and check if it is trivial.

If K is a sphere (PL or otherwise) then the fundamental group is trivial. From the affirmation of the PL Poincaré Conjecture by Smale [50] for dimensions $d \geq 5$ and by Perelman [44] for $d = 3$, we obtain the following converse.

Theorem 2 (PL Poincaré Conjecture) *Let K be a simply connected combinatorial d-manifold, $d \neq 4$, with the homology of the d-dimensional sphere. Then K is a PL sphere.*

Dimension $d = 4$ is a special case. Let M be a closed oriented combinatorial 4-manifold. This implies that M carries a unique smooth structure. We fix a generator $[M] \in H_4(M)$, which is called an *orientation class*. The restriction of the cohomology multiplication gives rise to the symmetric bilinear form

$$Q_M : H^2(M) \times H^2(M) \to \mathbb{Z} : (f, g) \mapsto \langle f \cup g, [M] \rangle$$

which is called the *intersection form* of M. Here $f \cup g$ denotes the cup product, and $\langle \cdot, \cdot \rangle$ is the evaluation of a cocycle on a cycle. Freedman proved that a simply connected 4-manifold with trivial intersection form is homeomorphic to the 4-sphere [28]. But his result does not say whether this also holds in the PL category. In fact, it is a major open problem whether or not exotic 4-spheres exist. Therefore, if Q_M is non-trivial this rules out the possibility that M is homeomorphic to $\mathbb{S}^4$, but the triviality of Q_M gives us nothing in the PL category since M could be an exotic sphere.

In all other dimensions $d \neq 4$, proving that $\pi_1(K) = 1$ is enough to recognize K to be a PL d-sphere (provided that K is a combinatorial d-manifold with spherical homology). In [49, Chapter 7] Seifert and Threlfall describe how one can obtain a finite presentation of $\pi_1(K)$ from any spanning tree in the 1-skeleton (with the remaining edges as generators) and all the 2-faces (as relators). However, checking if a finitely presented group is trivial is known to be undecidable. Discussing heuristic approaches to simplifying group presentations is beyond the scope of this paper. In practice we rely on `GAP` [30] which employs Tietze transformations.

(7) Perform random bistellar flips and check if the boundary of the $(d+1)$-simplex can be reached.

If all other tests are inconclusive, we can use a local search strategy to determine the PL type [12]. The boundary $\partial\Delta^{d+1}$ of the $(d+1)$-simplex is a d-dimensional simplicial complex with $d+2$ facets. A *bistellar flip* is a local modification of the combinatorial d-manifold K in which any subcomplex of K isomorphic to the star of a face in $\partial\Delta^{d+1}$ is replaced by its complementary facets. Two simplicial complexes are *bistellarly equivalent* if one is obtained from the other by a finite sequence of bistellar flips.

Theorem 3 (Pachner [43]) *A d-dimensional simplicial complex K is PL homeomorphic to the boundary of the $(d+1)$-simplex if and only if it is bistellarly equivalent to the boundary of the $(d+1)$-simplex.*

Algorithm 1 displays our strategy in a concise form.

Algorithm 1: Sphere recognition heuristics

Input: Hasse diagram of connected closed PL d-manifold K, where $d \geq 3$

Output: Decision: Is K PL homeomorphic to $\mathbb{S}^d$?

for N *rounds* **do**
  compute random acyclic matching
  if *this acyclic matching is spherical* **then return** YES
compute homology
if *homology not spherical* **then return** NO compute and simplify presentation of fundamental group π_1
if *this presentation is found to be trivial and* $d \neq 4$ **then return** YES **if** *this presentation is found to be non-trivial* **then return** NO
for N' *rounds* **do**
  perform random bistellar flip
  if *boundary of simplex reached* **then return** YES
return UNDECIDED

Notice that the ordering of the steps (4) through (7) is arbitrary. The specific ordering we chose here was useful for us to obtain the results from Section 4 below.

3 Data Structures and Implementation Details

The Algorithm 1 has been implemented in the mathematical software system `polymake` [31]. While the main focus is on convex polytopes, polyhedra and fans, `polymake` is also capable of dealing with simplicial complexes, matroids, graphs, tropical hypersurfaces and other objects. In the `polymake` project, `Perl` and `C++` are used as programming languages; our heuristics are implemented in `C++`. A description of some of the underlying data structures for Algorithm 1 can be found in the unpublished manuscript [35] and in [10, 12]. Here, we give a brief outline and some implementation details.

3.1 Hasse diagrams and face trees

The most obvious way to encode a finite abstract simplicial complex is to list its facets, each of which is written as a subset of the set of vertices. Depending on the kind of operations we intend to use, each facet can simply be encoded as an ordered list of vertices (e.g., implemented as `std::list`) or in a more involved fashion. Options for the latter include balanced search

trees, where the vertices of a given facet form the keys (e.g., implemented as `std::set` or `polymake::Set`), or as a bitstring, which writes a facet as a characteristic function on the set of vertices (e.g., implemented as `std::bitset` or `polymake::Bitset`). However, such a *facet list* encoding is not suitable for our heuristics.

A more sophisticated concept is provided by the *face trees* introduced by Kaibel and Pfetsch [38]; see also Ganter [29]. While this works in much greater generality, here we use it to construct the Hasse diagram of the face lattice of a simplicial complex K from the list of its facets. It is common to encode the entire Hasse diagram of K as a directed graph $H(K)$, where each node represents a face, and a directed arc from σ to τ means that the face σ is a maximal proper subface of τ. The unique source of this directed graph is the empty face. In what follows it is essential that each face is stored as an ordered sequence of vertices. This induces a lexicographic ordering of the faces of K of fixed dimension. Face trees are special in that for each face σ the data structure also explicitly stores the unique directed path from the empty face to σ which is lexicographically minimal. Constructing $H(K)$ using a breadth–first search yields an algorithm which is linear in the total number of faces of K.

Theorem 4 ([38, Theorem 5]) *The Hasse diagram of the face lattice of a simplicial complex K can be computed from the facet list of K in $O(d \cdot m^2 \cdot \phi)$ time, where m is the number of facets and ϕ is the total number of faces of$K$.*

Example 5 *Figure 1 shows the Hasse diagram of the standard triangulation of the real projective plane $\mathbb{RP}^2$ on six vertices. The face tree data structure encodes each face as a node together with a unique path from the bottom node to that face. For instance, associated with the face* 012 *is the path* $\emptyset \to 0 \to 01 \to 012$.

Face trees are implemented in `polymake` as the container type `polymake::FaceMap`. The term "map" is computer science lingo for "associative container". Here it is used as the type that comes with an interface to locate an arbitrary face of K in the graph $H(K)$ in $O(d \cdot m)$ time.

Remark 6 *The result in [38] is stated for face lattices of convex polytopes, but it also holds for arbitrary lattices which are both atomic and co-atomic. For simplicial complexes, the vertices are the atoms and the co-atoms are the facets. In [38] the authors use the parameter α to denote the number of vertex–facet incidences which, in the general poset picture, translates to the*

number of pairs (x, y) where x is an atom, y is a co-atom, and x is below y. For a pure d-dimensional simplicial complex, we have $\alpha = (d+1)m$ which is of order $O(d \cdot m)$.

The main advantage of using face trees over facet lists for our approach is that all face links are automatically represented as subtrees. This is immediately useful for our inductive approach as it avoids repeatedly checking the same face links (which may occur as the links of more than one face). More importantly, in Step (4) we construct the random acyclic matching in a destructive way. That is, we explicitly manipulate the arcs in the Hasse diagram. This enables the computation of data so large that they essentially take all the available main memory for just representing one Hasse diagram. The same holds for the bistellar flips in Step (7). These are also performed by explicitly changing the Hasse diagram locally. The face tree data structure is well-suited for deciding quickly if an intended bistellar flip is actually possible.

3.2 Homology computation

The Hasse diagram of a simplicial complex K describes all of its simplicial boundary matrices, and these matrices can be used for the homology computation. It is fairly easy to actually compute the homology of K if the coefficient domain is a field. In this case the boundary operator is a linear map. Its kernel and image are vector spaces. Hence the homology modules are vector spaces whose dimensions can be obtained as the differences between the co-rank and rank, respectively, of the two subsequent boundary matrices. Algorithmically, the rank (or co-rank) of a matrix can be obtained by a sequence of Gauss elimination steps.

Each integral homology module of K is a finitely generated abelian group. That is, they are a direct product of a finitely generated free abelian group and a product of finitely many cyclic groups of prime power order. On a conceptual level, the computation of integral homology differs only by a little from the computation of rational homology. There are two reasons for this: Firstly, $\mathbb{Q}$ is the quotient field of $\mathbb{Z}$. Secondly, $\mathbb{Z}$ is a Euclidean domain. The latter property comes in handy when one tries to perform a Gauss elimination over $\mathbb{Z}$. If the pivoting element is a unit, i.e., a 1 or -1, then the corresponding rational basis transformations are integral and the situation is then exactly the same as computing over $\mathbb{Q}$. Otherwise the pivoting element is not a unit, then Euclid's algorithm can be used to determine an integral transformation that will simplify the matrix. In this way,

after finitely many steps, each boundary matrix ∂_k can be transformed into a diagonal shape, the *Smith normal form* of ∂_k. It is then straightforward to determine the structure of $H_k(K)$ from the Smith normal forms of ∂_{k+1} and ∂_k. The existence and also the construction of the Smith normal form of an integral matrix is a classical result due to H. J. S. Smith [51]; for a modern account see Munkres [42, §11].

In terms of complexity there is a fundamental difference between Gauss elimination over $\mathbb{Q}$ and the sketched type of Euclidean-Gaussian elimination over $\mathbb{Z}$. In the rational case the computation time is bounded by a polynomial (in the size of the input matrix) since the growth of the coefficients can be controlled. This has been observed by Edmonds [24]. However, for the analogous operation over $\mathbb{Z}$ neither the size of the coefficients nor the number of arithmetic operations in $\mathbb{Z}$ is polynomially bounded. Kannan and Bachem [39] gave the first polynomial time Smith normal form algorithm, which was later improved by Iliopoulos [34] and others.

While the modular approach is valid for matrices with arbitrary integer coefficients, simplicial boundary matrices have entries 1, -1, and 0 only. That is to say, in an arbitrary simplicial boundary matrix it is *always* possible to perform at least a few Gauss elimination steps. Moreover, a typical boundary matrix is very sparse. If the matrix stays sparse during the elimination and if, additionally, one does not run out of unit coefficients too soon (such that it is possible to continue with elimination steps) an elimination based Smith normal form algorithm can be superior to the more sophisticated methods. This is why in practical applications elimination algorithms are often preferred. For a comprehensive survey, see Dumas et al [23].

Each edge in a matching in the Hasse diagram corresponds to a nonzero coefficient in some boundary matrix and vice versa. If the matching is acyclic a sequence of Gauss elimination steps can be performed so that the Gauss steps performed first do not destroy the (unit) pivots required for the subsequent steps.

Example 7 *Figure 2 shows the boundary matrices of $\mathbb{RP}^2_6$. The pivots corresponding to the acyclic matching from Figure 1 are marked. Using these pivots in an arbitrary order yields an elimination strategy for the computation of the homology modules.*

In Algorithm 1 the search for the acyclic matchings is listed before the homology computation. Clearly, this ordering can be reversed. More importantly, these steps can also be *intertwined* as finding an acyclic matching results in a strategy to compute the homology. To this end it is most natural to process the Hasse diagram from top to bottom level by level.

∂_2	01	02	03	04	05	12	13	14	15	23	24	25	34	35	45
012	1	−1				1									
014	1			-1				1							
023		1	−1							1					
035			1		−1									1	
045				1	-1										1
125						1			-1			1			
134							1	-1					1		
135							1		−1					1	
234										1	-1		1		
245											1	-1			1

$(\partial_1)^{\mathrm{tr}}$	01	02	03	04	05	12	13	14	15	23	24	25	34	35	45
0	-1	−1	−1	−1	−1										
1	1					-1	−1	−1	−1						
2		1				1				-1	−1	−1			
3			1				1			−1			-1	−1	
4				1				1			1		−1		1
5					1				1			1		1	1

Figure 2: Boundary matrices of $\mathbb{RP}^2_6$ with coefficients corresponding to acyclic matching marked.

This kind of homology computation is implemented, for instance, in `polymake` [31, 5], `CHomP` [20] and `RedHom` [17].

3.3 Random acyclic matchings

The `Random_Discrete_Morse` client implemented in `polymake` has three random strategies which we call *random-random*, *random-lex-first*, and *random-lex-last*. In this section, we will give a short outline of our implementation and describe the differences between the three strategies. The method has been described in [10].

Let K be a d-dimensional simplicial complex, which is not necessarily a manifold. A *free face* of K is a $(d-1)$-dimensional face that is contained in exactly one d-face. To save memory, our three strategies are destructive in the sense that they keep changing the complex K. In each step we try to pick one of the free faces and delete both it and the unique facet containing it from K. This is an elementary collapse, and the two removed faces form a regular pair, which is a matching edge in the Hasse diagram. The three strategies differ in how they pick the free face. If we run out of free faces we pick some facet, declare it critical and remove it. In both cases, after

removing a regular pair or after removing a critical face the dimension of the resulting complex, K', may drop to $d-1$. This process continues until K' is zero-dimensional. In this case K' only consists of vertices, all of which are declared critical.

For the random-random strategy, we first find *all* the free faces of K and collect them in a set or array data type. If this list is not empty, choose a free face uniformly at random. Taking the uniform distribution means that each free face has a fair chance of being taken, but this comes at a price since the sampling itself takes time if there are many free faces to choose from; see Knuth [40, §3.4.2]. If we run out of free faces the choice of the critical d-face is again uniformly at random.

Random-random is somehow the obvious strategy but there is a much cheaper way which maintains a certain amount of randomness. Here the price is that it seems to be next to impossible to say something about the resulting probability distribution. The idea is to randomly relabel the vertices of K once, at the beginning, and then to pick the free and critical faces in a deterministic way (which depends on the resulting labeling). Whenever a free or critical face is chosen, rather than selecting one at random, we pick the first (in the case of random-lex-first) or the last (in the case of random-lex-last) one.

The cost of being fair is quite significant when dealing with large complexes; for example, running the random-lex-first and random-lex-last strategies on `sd_4_bd_delta_4` took less than 3 minutes per run whereas the random-random strategy took approximately 2 hours per run.

The random-lex-last strategy was called "random-revlex" in [10]. We changed this here to random-lex-last to avoid confusion with (the term) reverse lexicographic ordering, which is different.

3.4 Random bistellar flips

In this section we discuss how randomness is applied in the algorithm of the `bistellar_simplification` client which was implemented in `polymake` by Nikolaus Witte and follows the simulated annealing strategy in [12]. The goal of the algorithm is to make local changes, called bistellar moves, to the input `SimplicialComplex` to lower its f-vector (lexicographically) as much as possible. Naturally, the algorithm prefers moves that lower the f-vector ("cooling"). Unfortunately, we may fall into a local minimum, i.e., when there are no moves to further improve the f-vector. At that point, we deliberately make moves that increase the f-vector for some number of rounds ("heating") then cool again, hoping that this will help jiggle us out

of that local minimum. The threshold for the number of rounds we tolerate before heating and the amount we heat is changed dynamically throughout the computation (greater when the complex is large and fewer when the complex is small).

Since each move only makes local changes, the algorithm only updates the list of possible moves instead of manipulating and carrying the entire Hasse diagram, which would be computationally very costly. More specifically, there is a subset `raw_options` of all the i-dimensional faces of the complex that are contained in exactly $d - i + 1$ facets, $0 \leq i \leq d$. The `raw_options` are updated after each move. But only some of the `raw_options` are 'proper' options, i.e., potential moves, which must satisfy the additional condition that performing the move does not introduce a face that is already present in the complex; see [12]. (The reason we do not update the list of proper options is because updating the raw options then checking is faster in the long run.)

The list of `raw_options` is grouped by dimension. The options that introduce new faces of lower dimension are heating moves and moves that introduce faces of higher dimension are cooling moves. During a cooling period we begin by checking for d-moves (where d is the dimension of the input complex) that remove a vertex from the complex. For this, start with a list of `raw_options` of dimension d. Create a random permutation of that list. Then we check each option in the permuted order until we find one that is proper. If none were proper, we move on to look for $(d-1)$-moves, and continue to dimension $d/2$ (or $(d+1)/2$ if d is odd) until we come across a proper option.

During a heating period, the story is slightly different. All raw options of dimension 0 are proper, i.e., all facets can be stellarly subdivided. For the heating strategy, the dimension of the heating move is chosen at random respecting a chosen distribution. For example, say we input a 4-dimensional complex. The default heat distribution for dimension 4 is `[10,10,1]`. So if the amount we heat `heating=210`, then the algorithm will heat for 210 rounds of which on average 10 are 0-moves, 100 are 1-moves and another 100 are 2-moves, but the order is chosen uniformly at random.

4 Examples and Experiments

The task of recognizing higher dimensional spheres or manifolds seems to be doomed given S. P. Novikov's non-recognizability result. Still, in many situations sphere recognition can be solved (easily) even for huge instances

Table 1: Collapsing the d-simplex.

d	rounds	non-perfect	percentage
8	10^9	12	0.0000012%
9	10^8	2	0.000002%
10	10^7	3	0.00003%
11	10^7	12	0.00012%
12	10^6	4	0.0004%
13	10^6	6	0.0006%
14	10^5	4	0.004%
15	10^5	8	0.008%
16	10^4	4	0.04%
17	10^4	10	0.1%
18	10^3	2	0.2%
19	10^3	6	0.6%
20	10^3	13	1.3%

with few exceptions. Recently, explicit triangulations of the Akbulut–Kirby 4-dimensional spheres [4] were constructed [55], for which the above recognition heuristics (4), (6), and (7) fail, but this outcome has thus far been unusual and isolated.

In this section, we delve into this discrepancy between the theoretical non-recognizability and the success we observed in practice in recognizing explicit examples. We will discuss limitations for the different recognition tests and experimentally determine a 'horizon' within which we can hope for effective recognition results.

Tests (1)–(3) of our integrated recognition approach always work and run in polynomial time; they were merely included to discard simplicial complexes that for obvious reasons are not manifolds.

4.1 First type of limitations for random discrete Morse test (4): Dimension of the input

Test (4) can be inconclusive. The main difficulty is that once we start collapsing a triangulated sphere (after the removal of an initial critical facet), we might encounter subcomplexes that are contractible, but non-collapsible. The most prominent example of a non-collapsible, contractible complex is the 2-dimensional *dunce hat* [58] which can be obtained from a single triangle by identifying, in a non-coherent way, its three boundary edges. The dunce hat can be triangulated as a simplicial complex with 8 vertices (see [8]),

Table 2: Spectrum for 10^9 runs on the 8-simplex.

discrete Morse vectors	count
(1 0 0 0 0 0 0 0 0)	999999988
(1 1 1 0 0 0 0 0 0)	4
(1 0 1 1 0 0 0 0 0)	7
(1 0 0 1 1 0 0 0 0)	1

while every contractible complex with fewer vertices is collapsible [6].

For Test (4), once an initial critical facet is removed from a manifold candidate, our goal is to show that the remaining simplicial complex collapses to a vertex and thus is a PL ball (see Theorem 1). The most basic example of a d-dimensional PL ball is a single d-dimensional simplex. Crowley et al showed that the 7-simplex with 8 vertices contains in its 2-skeleton an 8-vertex triangulation of the dunce hat onto which it collapses [21]. Our examples below are similar in spirit.

When we collapse the faces of a d-simplex using strategies random-lex-first or random-lex-last (see Section 3.3), we will always reach a single vertex (since a d-simplex is a cone and the respective collapses are towards an apex). However, if free faces are chosen randomly for a d-simplex, $d \geq 7$, we might run into the dunce hat or other contractible, but non-collapsible subcomplexes. For practical purposes, we immediately are interested in the rate of how often this is going to happen for the 7-simplex or higher-dimensional simplices. Table 1 displays the results of our random experiments. In dimension 7, all 10^9 rounds of sequences of random collapses were perfect. From dimension 8 on, we see a clear increase in the number of non-perfect discrete Morse vectors encountered, which means that randomly finding perfect discrete Morse vectors for the d-simplex beyond dimension 20 becomes increasingly more difficult. Although we have deterministic strategies to determine optimal discrete Morse vectors for various complexes, such as shellable complexes (which include single simplices), we cannot expect that running random collapses on a general 'random' input can perform better than on a single simplex.

Tables 2 and 3 give the actual discrete Morse vectors we found for the 8-simplex and the 20-simplex, respectively. We observe that we can get stuck (i.e., run out of free faces at a dimension $d > 0$) in different dimensions, as we already see for the 8-simplex in Table 2. While in the case of the 8-simplex we at most picked up two extra critical cells, the discrete Morse vector $(1, 0, 6, 48, 182, 377, 657, 876, 801, 493, 170, 22, 0, 0, 0, 0, 0, 0, 0, 0, 0)$ for

Table 3: Spectrum for 10^3 runs on the 20-simplex.

discrete Morse vectors	count
(1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	987
(1, 0, 0, 0, 6, 26, 59, 87, 61, 13, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 3, 30, 111, 158, 132, 82, 24, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 8, 34, 80, 126, 155, 126, 61, 27, 10, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 14, 27, 24, 13, 3, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 30, 117, 278, 409, 393, 213, 39, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 2, 25, 110, 236, 305, 175, 19, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 3, 5, 9, 34, 85, 134, 109, 33, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 19, 82, 150, 161, 90, 15, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 3, 18, 51, 118, 196, 264, 207, 57, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 11, 107, 243, 366, 463, 450, 261, 54, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 5, 30, 95, 160, 163, 124, 72, 27, 7, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 6, 48, 182, 377, 657, 876, 801, 493, 170, 22, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 0, 0, 0, 0, 0, 8, 14, 13, 14, 7, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1

the 20-simplex in Table 3 contains 3632 extra critical cells. Thus, in higher dimensions, not only do we get stuck with non-collapsible, contractible subcomplexes more often, but when we do get stuck, the resulting discrete Morse vectors will be huge. The first time we get stuck, the remaining complex is still homotopy equivalent to the d-simplex we started with. In particular, this reduced complex is *contractible*, but non-collapsible. This mean our heuristic procedure can also be used to search for such complexes.

In high dimensions, randomly finding spherical discrete Morse vectors for even the boundary d-sphere of a $(d+1)$-simplex is difficult, so testing combinatorial d-spheres that have more vertices and more facets will be even harder. Thus, a random search for spherical discrete Morse vectors will fail with high probability for *all* high-dimensional simplicial complexes, which renders Test (4) useless in high dimensions. Note that Tests (4) and (7) are most important in dimension 4; in all other dimensions they can, by the PL Poincaré conjecture, be replaced by the combination of Tests (5) and (6).

Let us return to the four examples of 2-dimensional contractible complexes we got stuck with when collapsing the 8-simplex; we call these four complexes D, S_a, S_b, and SQ. The first of these complexes D is a dunce hat as displayed in Figure 3. The other three complexes S_a, S_b, and SQ are not dunce hats, and they have inspired us to define a nice family of contractible, but non-collapsible 2-dimensional complexes that generalize the dunce hat.

Definition 8 *A* saw blade complex *is a 2-dimensional simplicial complex obtained from a triangulated disk by identifying vertices on the boundary of*

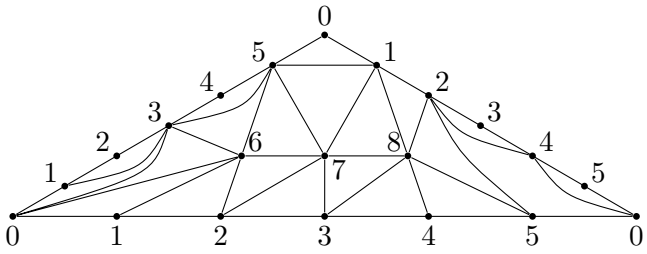

Figure 3: The dunce hat D.

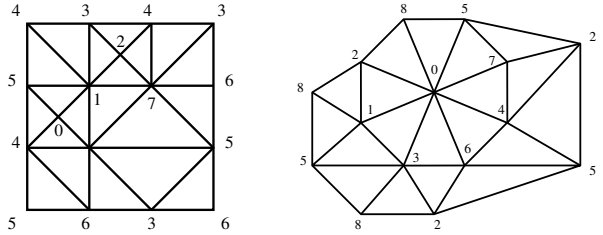

Figure 4: The two saw blade complexes S_a and S_b.

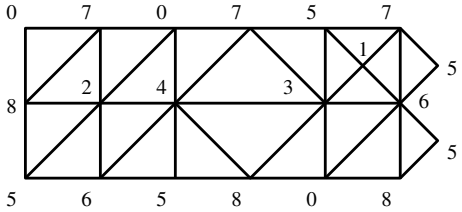

Figure 5: The complex SQ derived from the saw blade complex S.

the disk such that

- ▷ *the identification of the boundary vertices may not induce identifications of interior edges,*
- ▷ *the identification of the boundary vertices is* 3*-to-*1*, and*
- ▷ *the identification of the boundary edges is* 3*-to-*1*.*

And after identification

- ▷ *every boundary edge appears exactly twice with the same orientation and once with opposite orientation, and*
- ▷ *all identified boundary edges form a cycle.*

The two complexes S_a and S_b of Figure 4 are saw blade complexes; also every triangulation of the dunce hat is a saw blade complex.

In every saw blade complex the interior edges of the starting triangu-

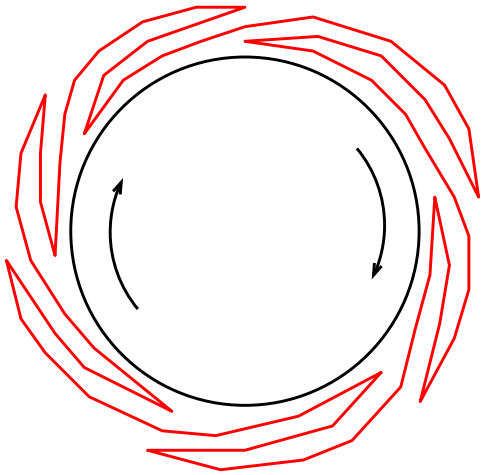

Figure 6: A saw blade labeling.

lated disk are contained in exactly two triangles, while the identified boundary edges lie in three triangles. We thus can uniquely (up to orientation) reconstruct from any saw blade triangulation the initial disk along with the boundary identifications. In S_a, the identified boundary edges form a 4-gon 4–3–6–5–4, and we have a three-fold folding of the 12-gon 4–3–4–3–6–5–6–3–6–5–4–5–4 onto the 4-gon 4–3–6–5–4.

Let us assume that the resulting cycle after identification is labeled 1–2–3–…–$(n-1)$–n–1. Along the identified boundary of the disk then every of the three copies of the vertex k is connected either to $k-1$ or $k+1$ (mod n), and the only way to arrange the vertices so that the conditions of Definition 8 are respected is in a *saw blade labeling* as indicated in Figure 6. We call the sequence of vertices starting at a turning point and then going forward, backward, and forward again till the next turning point is reached a *blade* of the saw blade labeling. A saw blade complex with k blades we call a *k-bladed saw blade complex*.

Theorem 9 *Let K be a saw blade complex. Then K is contractible, but non-collapsible.*

Proof. Any saw blade complex is non-collapsible. All edges either are of degree 2 or 3 so there are no free edges.

If K has one blade, K is a dunce hat and is contractible. Thus, let K have at least two blades. Let us choose two consecutive turning points and

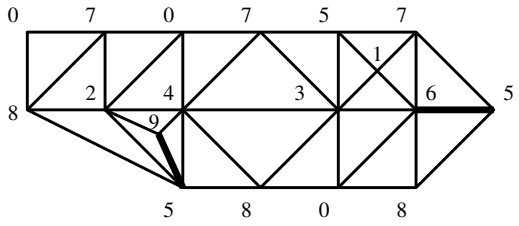

Figure 7: The saw blade complex S.

cut the identified disk into two parts P_1 and P_2 by inserting an interior arc from the one turning point to the next. As observed first in a special case by Hachimori [32], both P_1 and P_2 are contractible complexes (where the one retracts to the path of vertices of the blade between the consecutive turning points and the other retracts to the path formed by the remaining boundary vertices) that intersect in a contractible complex (the interior arc joining the two consecutive turning points). Thus, the union K of P_1 and P_2 is contractible. □

The fourth contractible, non-collapsible complex SQ (Figure 5) we found is not a saw blade complex, but can be obtained as a quotient from the saw blade complex S (Figure 7) by identifying the vertices 6 and 9. This way, the edge 5–6 becomes the same as the edge 5–9.

Proposition 10 *A k-bladed saw blade complex with $k \geq 3$ can be constructed with $3k$ vertices. A 2-bladed complex can be constructed with 9 vertices.*

Proof. For $k \geq 3$ we first give a construction with $4k$ vertices. Let the identified boundary of the saw blade disk be 1–2–1–2–3–2–3–4–. . . –$(n-1)$–n–1–n–1, i.e., we have k blades of length 1 each. In the interior of the disk we place a cycle with $3k$ vertices and connect the cycle vertices with the boundary cycle vertices in a zig-zag. Finally, the interior $3k$-gon can be triangulated without using additional vertices. At the turning points we can save an interior vertex each, yielding a construction with $3k$ vertices.

In the case of two blades, let the identified boundary be 1–2–1–2–3–1–3–2–3–1, where we need an extra vertex to avoid unwanted identifications. In the interior we then place a 6-gon and connect vertices similar to before.

In the case of the dunce hat with one blade, we need 8 vertices for a triangulation. □

Saw blade complexes with a different number of blades are combinatorially non-isomorphic complexes. Thus, as we see by the saw blade complexes

and their quotients, there is an abundance of 2-dimensional contractible, non-collapsible simplicial complexes on which we can get stuck when randomly searching for simplicial collapses — and there will be similar constructions in higher dimensions.

4.2 Second type of limitations for random discrete Morse test (4): Complicated triangulations

As pointed out in [2, 10], a random-lex-first or random-lex-last search for spherical discrete Morse vectors is often more promising than a general random search. However, there are triangulated spheres that do not have spherical discrete Morse vectors at all. The smallest known such example is an 18-vertex triangulation of the 3-dimensional sphere $\mathbb{S}^3$, constructed via a triple trefoil knot on three edges [9].

However, a knotted sphere is not what we usually have for the input. So what can we expect for some generic input? Standard sources for triangulated manifolds are examples from

- ▷ enumeration (where all enumeration schemes in the literature do not reach very far [15, 53]),
- ▷ combinatorial topology constructions such as suspensions, direct product triangulations, or connected sums (which are mostly harmless and do not give complicated triangulations),
- ▷ combinatorial constructions such as Hom complex constructions (which locally have a product like structure) [22] or moment-angle complexes [14].

Thus, most of the time we know the PL type of a triangulated manifold already by its construction and standard constructions usually result in easy-to-analyze triangulations as is also the case with applied data.

For one experiment, we started out with the boundary of the 4-simplex with 5 vertices, first added 525 vertices (via 0-moves in our bistellar flip implementation), then we performed 50,000 1-moves at random, before we run further 10^6 rounds of bistellar flips where we allowed both random 1- and 2-moves. For the resulting 'random' triangulation of the 3-sphere with face vector $f = (530, 50474, 99888, 49944)$, we computed discrete Morse vectors. The smallest vector found was $(1, 2192, 2192, 1)$ — far away from the perfect vector $(1, 0, 0, 1)$. But, using bistellar flips again, we were easily able to reduce the complex back to the boundary of the simplex. We also used `GAP` to find a trivial presentation for the fundamental group of the example, which took 16 hours for the simplification.

In an effort to obtain small, but highly non-trivial triangulations of the 4-dimensional sphere, explicit triangulations of the Akbulut–Kirby 4-

spheres [4] were constructed in [55]. This series of spheres is based on non-trivial presentations of the trivial group (as the fundamental group of the examples) that cannot be easily transformed into trivial presentations. The resulting spheres AK(r), $r \geq 3$, are standard PL spheres [3] and have the non-trivial presentations built into their 2-skeletons and thus forcing the fundamental group to become hard to recognize. Test (4) fails on all examples of this series, Test (7) worked for $r = 3$, but failed for $r \geq 4$. For $r = 4$, Test (6) was successful in some runs, however, in dimension 4 this only determines (in combination with Test (5)) that the input is a topological 4-sphere, but yielding no information on the PL type. For $r \geq 5$, none of the tests (4), (6), (7) produced positive results.

4.3 Third type of limitations for random discrete Morse test (4): Number of vertices or facets

As already pointed out, it is often rather easy to find optimal discrete Morse vectors even for huge (nicely structured) complexes; see [10]. However, for higher-dimensional simplices or for higher barycentric subdivisions in fixed dimension $d \geq 4$ of boundaries of $(d+1)$-simplices [2] we might get stuck in substructures like the dunce hat. The probability to indeed encounter a dunce hat or a similar contractible, but non-collapsible subcomplex is extremely small in low dimension, as we have seen in our experiments above.

As mentioned earlier, Adiprasito showed that a sufficiently large iterated barycentric subdivision of any PL sphere admits a spherical discrete Morse function. Yet, the average number of critical cells for random discrete Morse vectors grows exponentially with the number of barycentric subdivisions [2]. We ran our implementation on higher barycentric subdivisions of boundaries of simplices. For the 3rd barycentric subdivision `sd_3_bd_delta_4` of the boundary of the 4-simplex with $f = (12600, 81720, 138240, 69120)$ the optimal discrete Morse vector $(1, 0, 0, 1)$ was found in 994 out of 1000 runs of the random-lex-last version [2] of the random discrete Morse search; see Table 4. For the 4-th barycentric subdivision `sd_4_bd_delta_4` of the boundary of the 4-simplex with face vector $f = (301680, 1960560, 3317760, 1658880)$ the optimal discrete Morse vector $(1, 0, 0, 1)$ was found in only 844 out of 1000 runs, which may indicate that the horizon for computations lies near the 5-th barycentric subdivision.

Table 4: Distribution of discrete Morse vectors for barycentric subdivisions.

random-random		**random-lex-first**		**random-lex-last**	
`sd_3_bd_delta_4` with $f = (12600, 81720, 138240, 69120)$					
(1,0,0,1):	1000	**(1,0,0,1)**:	999	**(1,0,0,1)**:	994
		$(1,1,1,1)$:	1	$(1,1,1,1)$:	6
`sd_4_bd_delta_4` with $f = (301680, 1960560, 3317760, 1658880)$					
(1,0,0,1):	20	**(1,0,0,1)**:	829	**(1,0,0,1)**:	844
		$(1,1,1,1)$:	143	$(1,1,1,1)$:	107
		$(1,2,2,1)$:	19	$(1,2,2,1)$:	30
		$(2,3,2,1)$:	3	$(1,3,3,1)$:	9
		$(2,5,4,1)$:	2	$(1,4,4,1)$:	4
		$(1,3,3,1)$:	2	$(2,5,4,1)$:	2
		$(1,4,4,1)$:	1	$(1,5,5,1)$:	2
		$(1,5,5,1)$:	1	$(2,3,2,1)$:	1
				$(2,7,6,1)$:	1

4.4 Limitations for homology test (5)

In this work we have not evaluated the performance of homology computations. But since standard homology packages such as `polymake` [31, 5], `CHomP` [20], `RedHom` [17], and `Perseus` [45] rely on a discrete Morse type reduction (where it is NP-hard to find an optimal topological simplification) as a preprocessing step for (polynomial time) Smith normal form computations, we should expect to see similar effects as we have experienced for Test (4). To be precise, we might encounter a 'horizon' for the dimension and complex size (in terms of vertices and facets), after which 'standard' input can no longer be processed effectively. That is, our experimentation suggests that the advantage gained from the discrete Morse preprocessing will probably diminish for very large complexes.

4.5 Limitations for fundamental group test (6)

It seems to be particularly difficult to construct a triangulation of a simply connected manifold for which Test (6) fails. One interesting class of examples are provided by the explicit triangulations of the Akbulut–Kirby 4-spheres in [55].

The triangulations AK(r) of the Akbulut–Kirby 4-spheres are defined via the (non-trivial) presentations $G(r) = \langle\, x, y \mid xyx = yxy;\, x^r = y^{r-1} \,\rangle$ of

the trivial group and have face vector $f = (176 + 64r, 2390 + 1120r, 7820 + 3840r, 9340 + 4640r, 3736 + 1856r)$ for $r \geq 3$. For $r = 3$ only, we were able to use bistellar flips to reduce AK(3) to the boundary of the 5-simplex. For $r = 4, 5$, running bistellar flips eventually resulted in small triangulations with about 30 vertices. Test (4) failed on all of these triangulations. For $r = 4$, we started bistellar flips on AK(4) for different random seeds and then ran Test (6) on the small triangulations we obtained. In 100 out of 450 runs a trivial presentation was found for the fundamental group, giving a certificate that AK(4) is a topological 4-sphere (by Freedman's classification), but failing to give a certificate for whether AK(4) is a PL 4-sphere (which a priori is known by Akbulut's proof [3] that all the examples AK(r) are standard). For $r = 5$, all three Tests (4), (6), and (7) failed, but for one of the runs for $r = 5$ and every run for $r = 3$ we actually obtained, again, the initial presentations $G(r)$ *after* simplification of the presentations with `GAP`— which renders the series AK(r) as an interesting testing ground for sphere recognition heuristics.

In dimension 3, the known recognition algorithms for the 3-sphere make use of normal surface theory. As a byproduct of Perelman's proof of the 3-dimensional Poincaré conjecture, 3-sphere recognition algorithms can in fact be used to recognize whether the fundamental group of a triangulated 3-manifold is trivial — the only such 3-manifold is the 3-sphere.

4.6 Limitations for bistellar simplification test (7)

The Akbulut–Kirby 4-spheres AK(r) are the only explicit examples of combinatorial spheres we know of for which we have not been able to find a simplification using bistellar flips. In higher dimensions choosing appropriate schemes for the flip options is non-trivial, which may result in a failure of Test (7).

4.7 A contractible non-5-ball.

Finding interesting and challenging test examples for our recognition procedure is non-trivial. Most examples from the literature are tiny, can easily fit into memory, and can be recognized instantaneously. A recent example of larger size is `contractible_non_5_ball` [2], a non-PL triangulation of a contractible and collapsible 5-manifold, different from the 5-ball, with $f = (5013, 72300, 290944, 495912, 383136, 110880)$. The example was shown to be collapsible [2]; the vector $(1, 0, 0, 0, 0, 0)$ was obtained after only a single random discrete Morse vector search in a running time of 82 hours with a

Table 5: Distribution of discrete Morse vectors for a contractible non-5-ball.

random-random		**random-lex-first**		**random-lex-last**	
`contractible_non_5_ball`					
(1,0,0,0,0,0):	100	**(1,0,0,0,0,0)**:	967	**(1,0,0,0,0,0)**:	1000
		$(1,1,1,0,0,0)$:	17		
		$(1,2,2,0,0,0)$:	4		
		$(1,3,3,0,0,0)$:	3		
		$(1,0,1,1,0,0)$:	3		
		$(1,1,3,2,0,0)$:	2		
		$(1,0,2,2,0,0)$:	2		
		$(1,0,3,3,0,0)$:	1		
		$(1,1,4,3,0,0)$:	1		

`GAP` implementation. Our new implementation in the `polymake` system [31] of the search for random discrete Morse vectors produced the same result in about 9 seconds with the `random-lex-first` and `random-lex-last` strategies and in about 10 minutes with the `random-random` strategy; the computations ran on a standard desktop computer with AMD Phenom II X6 1090T CPU (3.2 GHz, 6422 bogomips) and 8 GB RAM; see Table 5 for a distribution of discrete Morse vectors we found for the example for the different strategies `random-random`, `random-lex-first`, and `random-lex-last`.

The boundary of `contractible_non_5_ball` is a combinatorial 4-manifold `contractible_non_5_ball_boundary` with face vector $f = (5010, 65520, 212000, 252480, 100992)$. We used Test (4) to confirm that this example is indeed a combinatorial manifold. In fact, for every face link a spherical discrete Morse vector was found in a single try. In total, the recognition of all face links took about 7.5 hours. The example `contractible_non_5_ball_boundary` is a homology 4-sphere that has the binary icosahedral group as its fundamental group, as was confirmed computationally in [2].

5 Conclusion

Sphere recognition — despite its theoretical constraints — has proven to be often feasible in practice, even in high dimensions. Our interest in this work is mainly of a topological nature. But with the recent surge in development of topological methods to analyze large data sets, the complicated examples for which our heuristics fail may become useful for testing and improving those methods.

Acknowledgements

We are grateful to Bruno Benedetti, Vidit Nanda, and Konstantin Mischaikow for helpful discussions.

References

References

[1] K. Adiprasito and B. Benedetti. Subdivisions, shellability, and collapsibility of products. arXiv:1202.6606, 2012, 17 pages; *Combinatorica* (to appear).

[2] K. Adiprasito, B. Benedetti, and F. H. Lutz. Extremal examples of collapsible complexes and random discrete Morse theory. arXiv:1404.4239, 2014, 20 pages.

[3] S. Akbulut. Cappell–Shaneson homotopy spheres are standard. *Ann. Math.*, 171:2171–2175, 2010.

[4] S. Akbulut and R. Kirby. A potential smooth counterexample in dimension 4 to the Poincaré conjecture, the Schoenflies conjecture, and the Andrews–Curtis conjecture. *Topology*, 24:375–390, 1985.

[5] Benjamin Assarf, Ewgenij Gawrilow, Katrin Herr, Michael Joswig, Benjamin Lorenz, Andreas Paffenholz, and Thomas Rehn. `polymake` in linear and integer programming, 2014. preprint arXiv:1408.4653.

[6] B. Bagchi and B. Datta. Combinatorial triangulations of homology spheres. *Discrete Math.*, 305:1–17, 2005.

[7] B. Benedetti. Smoothing discrete Morse theory. arXiv:1212.0885, 2014, 27 pages; *Annali della Scuola Normale Superiore di Pisa Classe di Scienze* (to appear).

[8] B. Benedetti and F. H. Lutz. The dunce hat and a minimal non-extendably collapsible 3-ball. Electronic Geometry Model, No. 2013.10.001, 2013. `http://www.eg-models.de/2013.10.001`.

[9] B. Benedetti and F. H. Lutz. Knots in collapsible and non-collapsible balls. *Electron. J. Comb.*, 20(3):P31, 2013.

[10] B. Benedetti and F. H. Lutz. Random discrete Morse theory and a new library of triangulations. *Exp. Math.*, 23:66–94, 2014.

[11] B. Benedetti and G. M. Ziegler. On locally constructible spheres and balls. *Acta Math.*, 206:205–243, 2011.

[12] A. Björner and F. H. Lutz. Simplicial manifolds, bistellar flips and a 16-vertex triangulation of the Poincaré homology 3-sphere. *Exp. Math.*, 9:275–289, 2000.

[13] U. Brehm and W. Kühnel. 15-vertex triangulations of an 8-manifold. *Math. Ann.*, 294:167–193, 1992.

[14] V. M. Buchstaber and T. E. Panov. Moment-angle complexes and combinatorics of simplicial manifolds. *Russ. Math. Surv.*, 55:567–569, 2000. Translation from *Usp.Mat.Nauk*, 55(3):171–172, 2000; extended version arXiv:math/0005199.

[15] B. A. Burton. Simplification paths in the Pachner graphs of closed orientable 3-manifold triangulations. arXiv:1110.6080, 2011, 39 pages.

[16] J. W. Cannon. Shrinking cell-like decompositions of manifolds. Codimension three. *Ann. Math.*, 110:83–112, 1979.

[17] CAPD::RedHom. Redhom software library. `http://redhom.ii.uj.edu.pl`.

[18] M. K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Math.*, 217(1-3):101–113, 2000. Formal Power Series and Algebraic Combinatorics (Vienna, 1997).

[19] A. V. Chernavsky and V. P. Leksine. Unrecognizability of manifolds. *Ann. Pure Appl. Logic*, 141:325–335, 2006.

[20] CHomP. *C*omputational *H*omology *P*roject. `http://chomp.rutgers.edu`.

[21] K. Crowley, A. Ebin, H. Kahn, P. Reyfman, J. White, and M. Xue. Collapsing a simplex to a noncollapsible simplicial complex. Preprint, 2003, 7 pages.

[22] P. Csorba and F. H. Lutz. Graph coloring manifolds. In C. A. Athanasiadis, V. V. Batyrev, D. I. Dais, M. Henk, and F. Santos, editors, *Algebraic and Geometric Combinatorics,* Euroconf. Math., Algebraic and Geometric Combinatorics, Anogia, Crete, Greece, 2005, volume 423 of *Contemporary Mathematics*, pages 51–69, Providence, RI, 2006. American Mathematical Society.

[23] J. G. Dumas, F. Heckenbach, D. Saunders, and V. Welker. Computing simplicial homology based on efficient smith normal form algorithms. In M. Joswig and N. Takayama, editors, *Algebra, Geometry, and Software Systems*, pages 177–206. Springer, Berlin, 2003.

[24] J. Edmonds. Systems of distinct representatives and linear algebra. *J. Res. Nat. Bur. Standards Sect. B*, 71B:241–245, 1967.

[25] R. D. Edwards. The double suspension of a certain homology 3-sphere is S^5. *Notices AMS*, 22:A–334, 1975.

[26] R. Forman. Morse theory for cell complexes. *Adv. Math.*, 134:90–145, 1998.

[27] R. Forman. A user's guide to discrete Morse theory. *Sémin. Lothar. Comb.*, 48:B48c, 35 p., electronic only, 2002.

[28] Michael Hartley Freedman. The topology of four-dimensional manifolds. *J. Differ. Geom.*, 17(3):357–453, 1982.

[29] B. Ganter. Algorithmen zur formalen Begriffsanalyse. In *Beiträge zur Begriffsanalyse (Darmstadt, 1986)*, pages 241–254. Bibliographisches Inst., Mannheim, 1987.

[30] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.7.6*, 2014.

[31] E. Gawrilow and M. Joswig. polymake: a framework for analyzing convex polytopes. In *Polytopes—combinatorics and computation (Oberwolfach, 1997)*, volume 29 of *DMV Sem.*, pages 43–73. Birkhäuser, Basel, 2000.

[32] M. Hachimori. Decompositions of two-dimensional simplicial complexes. *Discrete Math.*, 308:2307–2312, 2008.

[33] J. Hass and G. Kuperberg. The complexity of recognizing the 3-sphere. *Oberwolfach Reports*, 24:1425–1426, 2012.

[34] C. S. Iliopoulos. Worst-case complexity bounds on algorithms for computing the canonical structure of finite abelian groups and the Hermite and Smith normal forms of an integer matrix. *SIAM J. Comput.*, 18(4):658–669, 1989.

[35] M. Joswig. Computing invariants of simplicial manifolds. arXiv:math/0401176, 2004, 13 pages.

[36] M. Joswig, F. H. Lutz, and M. Tsuruga. Heuristics for sphere recognition. In H. Hong and C. Yap, editors, *Mathematical Software — ICMS 2014*, volume 8592 of *LNCS*, pages 152–159. Springer, 2014.

[37] M. Joswig and M. E. Pfetsch. Computing optimal Morse matchings. *SIAM J. Discrete Math.*, 20:11–25, 2006.

[38] V. Kaibel and M. E. Pfetsch. Computing the face lattice of a polytope from its vertex-facet incidences. *Comput. Geom.*, 23:281–290, 2002.

[39] R. Kannan and A. Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. *SIAM J. Comput.*, 8(4):499–507, 1979.

[40] D. E. Knuth. *The Art of Computer Programming. Vol. 2.* Addison-Wesley, Reading, MA, 1998. Seminumerical algorithms, Third edition [of MR0286318].

[41] T. Lewiner, H. Lopes, and G. Tavares. Optimal discrete Morse functions for 2-manifolds. *Comput. Geom.*, 26:221–233, 2003.

[42] J. R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley Publishing Company, Menlo Park, CA, 1984.

[43] U. Pachner. Konstruktionsmethoden und das kombinatorische Homöomorphieproblem für Triangulationen kompakter semilinearer Mannigfaltigkeiten. *Abh. Math. Sem. Univ. Hamburg*, 57:69–86, 1987.

[44] G. Perelman. The entropy formula for the Ricci flow and its geometric applications. arXiv:math/0211159, 2002, 39 pages.

[45] Perseus. `http://www.math.rutgers.edu/~vidit/perseus.html`.

[46] J. H. Rubinstein. An algorithm to recognize the 3-sphere. In *Proc. of the ICM 1(2), Zürich, 1994*, pages 601–611. Birkhäuser, Basel, 1995.

[47] E. Saucan, E. Appleboim, and Y. Y. Zeevi. Sampling and reconstruction of surfaces and higher dimensional manifolds. *J. Math. Imaging Vision*, 30(1):105–123, 2008.

[48] S. Schleimer. Sphere recognition lies in NP. In *Low-dimensional and symplectic topology*, volume 82 of *Proc. Sympos. Pure Math.*, pages 183–213. Amer. Math. Soc., Providence, RI, 2011.

[49] H. Seifert and W. Threlfall. *Lehrbuch der Topologie*. B. G. Teubner, Leipzig, 1934.

[50] S. Smale. Generalized Poincaré's conjecture in dimensions greater than four. *Ann. Math.*, 74:391–406, 1961.

[51] H. J. S. Smith. On systems of linear indeterminate equations and congruences. *Philos. Trans. Roy. Soc. London Ser. A*, 151:293–326, 1861.

[52] J. Spreer and W. Kühnel. Combinatorial properties of the $K3$ surface: Simplicial blowups and slicings. *Exp. Math.*, 20:201–216, 2011.

[53] T. Sulanke and F. H. Lutz. Isomorphism free lexicographic enumeration of triangulated surfaces and 3-manifolds. *Eur. J. Comb.*, 30:1965–1979, 2009.

[54] A. Thompson. Thin position and the recognition problem for S^3. *Math. Res. Lett.*, 1(5):613–630, 1994.

[55] M. Tsuruga and F. H. Lutz. Constructing complicated spheres. In *EuroCG 2013*, pages 29–32, 2013.

[56] I. A. Volodin, V. E. Kuznetsov, and A. T. Fomenko. The problem of discriminating algorithmically the standard three-dimensional sphere. *Russ. Math. Surv.*, 29(5):71–172, 1974.

[57] J. H. C. Whitehead. Simplicial spaces, nuclei and m-groups. *Proc. Lond. Math. Soc., II. Ser.*, 45:243–327, 1939.

[58] E. C. Zeeman. On the dunce hat. *Topology*, 2:341–358, 1963.